

Accepted manuscript

Moreno, E. R. & Beferull-Lozano, B. (2023). An Online Multiple Kernel Parallelizable Learning Scheme. IEEE Signal Processing Letters, 31, 121-125.
<https://doi.org/10.1109/LSP.2023.3343185>.

Submitted to: IEEE Signal Processing Letters

DOI: <https://doi.org/10.1109/LSP.2023.3343185>

AURA: <https://hdl.handle.net/11250/3126155>

Copyright: © 2023 IEEE

License: Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

An Online Multiple Kernel Parallelizable Learning Scheme

Emilio Ruiz-Moreno* and Baltasar Beferull-Lozano*[†], *Senior Member, IEEE*

*WISNET Center, Department of ICT, University of Agder, Grimstad, Norway

[†]SIGIPRO Department, Simula Metropolitan Center for Digital Engineering, Oslo, Norway

Abstract—The performance of reproducing kernel Hilbert space-based methods is known to be sensitive to the choice of the reproducing kernel. Choosing an adequate reproducing kernel can be challenging and computationally demanding, especially in data-rich tasks without prior information about the solution domain. In this paper, we propose a learning scheme that scalably combines several single kernel-based online methods to reduce the kernel-selection bias. The proposed learning scheme applies to any task formulated as a regularized empirical risk minimization convex problem. More specifically, our learning scheme is based on a multi-kernel learning formulation that can be applied to widen any single-kernel solution space, thus increasing the possibility of finding higher-performance solutions. In addition, it is parallelizable, allowing for the distribution of the computational load across different computing units. We show experimentally that the proposed learning scheme outperforms the combined single-kernel online methods separately in terms of the cumulative regularized least squares cost metric.

Index Terms—Online learning, reproducing kernel Hilbert space, multi-kernel learning.

I. INTRODUCTION

Reproducing kernel Hilbert space (RKHS)-based methods allow modeling highly non-linear relationships at a moderate computational cost [1]. Thanks to their simplicity and generality, they have been successfully adopted in a wide range of signal-processing applications [2], [3].

The performance of any RKHS-based method strongly relies on a preselected reproducing kernel (RK). The efficient selection of an adequate RK presumes some task-specific prior information, such as knowledge about the data domain, invariant data transformations, geometrical data structures, or some properties of the underlying data generating process [4]. For example, spline interpolation RKs are best suited for smooth data [5], [6]. Similarly, radial basis RKs can perform poorly if their associated hyperparameters are not properly tuned to the task. The kernel-selection issue cannot be easily mitigated via cross-validation [7] because the associated computational load grows prohibitively with the number of RKs. On the other hand, efficiently computable approximations of the leave-one-out error [8] or hyperparameter optimization techniques [9] usually involve non-convexity and may lead to undesirable local minima.

This work was supported by the SFI Offshore Mechatronics grant 237896/O30 and the PETROMAKS Smart-Rig grant 244205/E30 from the Research Council of Norway.

Multi-kernel methods compensate for the lack of task-specific prior information using a predefined set of RKs known as *dictionary*. The dictionary can be formed by integrating different types of RKs, the same RK with different hyperparameter values, or a mix of both. Typically, the preselected RK is constructed as a combination of several RKs from the dictionary [10]. Therefore, how the dictionary is formed and how the preselected RK is constructed have a pivotal impact on the resulting accuracy and complexity of the method. For instance, the larger the dictionary is, the more likely it is to reduce the kernel-selection bias compared to a particular RK or hyperparameter choice. In addition, larger dictionaries allow for greater adaptability when learning from data samples since, in practice, these samples may come from a combination of different sources. On the other hand, increasing the dictionary size becomes computationally demanding or even prohibitive. For this reason, a commonly sought goal for multi-kernel methods is to find a compromise between performance and a computationally light and compact representation of the proposed solution in terms of the dictionary elements [11].

Another related scaling issue that also applies to single-kernel methods is the curse of kernelization, i.e., potentially unbounded linear growth in model size with the amount of data [12]. This drawback is generally addressed through online approaches, which may rely on sparsification procedures [13], [14], [15], [16], [17] or dimensionality reduction approximations [18], [19], among others [20].

In this context, some works [21], [22], [23] have explored the use of online methods for determining the optimal solution associated with each single RK within the dictionary, as well as the best combination of these single-kernel solutions under a given task. Following a conceptually similar approach, this paper proposes a novel multi-kernel learning scheme that can be parallelized across RKs by efficiently combining the solution of several single RK-based online methods running concurrently. This provides scalability with respect to the number of data samples and adaptability across different data patterns. Moreover, it allows to distribute the computational load across different computing units as the dictionary size increases. Our proposed scheme applies to any task that can be formulated as a regularized empirical risk minimization (RERM) convex problem [24], [25]. Finally, the performance of the proposed learning scheme is experimentally validated in terms of the cumulative regularized least squares cost metric.

II. PROBLEM FORMULATION

Supervised learning is arguably one of the core topics in machine learning [26]. Many supervised learning tasks can be formulated as RERM convex problems whose solution admits a kernel representation. That is, given a set of N data samples $\mathcal{S} = \{(x^{(n)}, y^{(n)})\}_{n=1}^N \subseteq \mathcal{X} \times \mathcal{Y}$, and an RKHS¹ $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, the goal is to find a function estimate $f \in \mathcal{H}$ minimizing the following regularized functional cost

$$\mathcal{C}_\eta(f; \mathcal{S}) = \sum_{n=1}^N \ell(f(x^{(n)}), y^{(n)}) + \frac{\eta}{2} \|f\|_{\mathcal{H}}^2, \quad (1)$$

where the loss $\ell: \mathcal{Y}^2 \rightarrow \mathbb{R} \cup \{\infty\}$ is a proper convex function used as a goodness-of-fit metric, the regularizer $\|\cdot\|_{\mathcal{H}}^2: \mathcal{H} \rightarrow \mathbb{R}$ is the squared RKHS \mathcal{H} induced norm, and the hyperparameter $\eta \in \mathbb{R}_+$ controls the model complexity of the solution.

Under a multi-kernel learning framework, one typically constructs a valid RK by adequately combining the RKs within a preselected dictionary [10]. Particularly, finding the RK within a convex hull of P positive definite RKs that yields the function estimate incurring the lowest functional cost (1) is equivalent to obtaining a solution from \mathcal{H} built as the RKHS direct sum $\mathcal{H}_1 \oplus \dots \oplus \mathcal{H}_P$, where each p th RKHS $\mathcal{H}_p = \overline{\text{span}}\{k_p(x, \cdot) : x \in \mathcal{X}\}$, being $k_p: \mathcal{X}^2 \rightarrow \mathbb{R}$ its associated RK [27]. The solution $f \in \mathcal{H}$ that minimizes (1), can be expressed without loss of generality as $f = \boldsymbol{\theta}^\top \mathbf{f}$, where $\mathbf{f} = [f_1, \dots, f_P]^\top \in \mathcal{H}_{1:P}$, with $\mathcal{H}_{1:P} \triangleq \mathcal{H}_1 \times \dots \times \mathcal{H}_P$ and $\boldsymbol{\theta} = [\theta_1, \dots, \theta_P]^\top \in \Delta^{P-1}$, with $\Delta^{P-1} \triangleq \{\boldsymbol{\beta} \in \mathbb{R}^P : \boldsymbol{\beta} \succeq \mathbf{0} \text{ and } \mathbf{1}^\top \boldsymbol{\beta} = 1\}$ denoting a simplex [28]. Thus, the RERM problem posed before becomes

$$\min_{\boldsymbol{\theta} \in \Delta^{P-1}, \mathbf{f} \in \mathcal{H}_{1:P}} \mathcal{C}_\eta(f; \mathcal{S}) \quad (2a)$$

$$\text{subject to: } f = \boldsymbol{\theta}^\top \mathbf{f}. \quad (2b)$$

Optimization problem (2) is bi-convex, meaning that it is convex in $\boldsymbol{\theta}$ for a fixed \mathbf{f} and vice-versa. Still, it is not jointly convex in both optimization variables. It can be tackled via specialized methods that primarily exploit the convex substructures of the problem [29]. However, these methods do not scale well with the number of RKs and data samples, denoted as P and N , respectively.

This paper presents a method to solve efficiently (2) for large P and N .

III. PROPOSED SOLUTION

This section describes how to synergize an online formulation and an upper bound on the objective (2a) to solve (2) scalably with respect to N and P .

A. Online setting

Online settings [30], [31] can be adopted to solve (2) achieving low run-time complexity with respect to N while incurring a certain tolerable (cumulative) cost. They usually trade-off solution accuracy for speed, e.g., by processing only a few samples every iteration, for low memory complexity,

¹The notation $\mathcal{Y}^{\mathcal{X}}$ refers to the set of functions from \mathcal{X} to \mathcal{Y} .

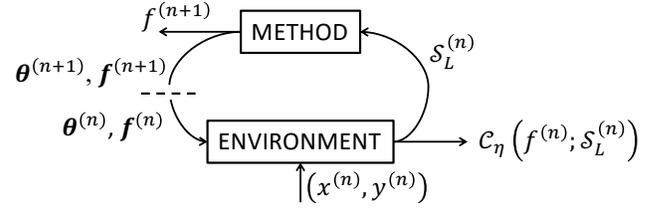


Fig. 1. Visual description of the online setting considered.

e.g., by discarding samples after a few processing steps, or for model complexity control, i.e., bounded model size regardless of whether N increases.

The online setting considered in this work can be cast as a method-environment iterative game [32]. The data samples in \mathcal{S} are assumed to be available sequentially. Then, at each iteration step n , the method chooses a function estimate $f^{(n)} \in \mathcal{H}_1 \oplus \dots \oplus \mathcal{H}_P$ expressible as $f^{(n)} = \boldsymbol{\theta}^{(n)\top} \mathbf{f}^{(n)}$ with $\boldsymbol{\theta}^{(n)} \in \Delta^{P-1}$ and $\mathbf{f}^{(n)} \in \mathcal{H}_{1:P}$. In response, the environment penalizes the proposed function estimate $f^{(n)}$ with the incurred cost $\mathcal{C}_\eta(f^{(n)}; \mathcal{S}_L^{(n)})$, where $\mathcal{S}_L^{(n)} = \{(x^{(i)}, y^{(i)})\}_{i=n-L}^n \subseteq \mathcal{S}$ is a sliding window of L data samples and $n_L \triangleq \max\{n-L+1, 1\}$. Finally, once the n th function estimate $f^{(n)}$ is chosen, the method receives the n th data window $\mathcal{S}_L^{(n)}$, which can be used at the next iteration² step $n+1$. Fig. 1 visually describes the procedure.

B. Upper bound on the functional cost

As we show next, we can improve scalability with respect to the number of RKs, by making use of the following upper bound:

$$\mathcal{C}_\eta(f; \mathcal{S}) = \sum_{n=1}^N \ell(\boldsymbol{\theta}^\top \mathbf{f}(x^{(n)}), y^{(n)}) + \frac{\eta}{2} \|\boldsymbol{\theta}^\top \mathbf{f}\|_{\mathcal{H}}^2 \quad (3a)$$

$$\leq \sum_{n=1}^N \sum_{p=1}^P \theta_p \ell(f_p(x^{(n)}), y^{(n)}) + \theta_p^2 \frac{\eta}{2} \|f_p\|_{\mathcal{H}_p}^2 \quad (3b)$$

$$\triangleq \check{\mathcal{C}}_\eta(\boldsymbol{\theta}, \mathbf{f}; \mathcal{S}). \quad (3c)$$

The first upper-bounded term in (3b) follows directly from Jensen's inequality [33], whereas the second term is obtained by invoking the definition of the RKHS direct sum norm [34]. Even though the second term in (3b) could have also been upper bounded through Jensen's inequality, as in [28], exploiting the definition of the RKHS direct sum norm, leads to a tighter bound because $\theta_p^2 \leq \theta_p$ for $\theta_p \in [0, 1]$.

The key advantage of the upper bound cost (3c) is that it is separable across the P RKs within the dictionary, hence allowing for parallelization at the expense of some loss in optimality, albeit with still satisfactory performance.

C. Parallelizable learning scheme

Our proposed learning scheme consists of executing at each iteration step n the following consecutive operations:

²At the first iteration step $n=1$, the method has not received any data sample, thus $f^{(1)}$ is set as some arbitrary initial function estimate.

Algorithm 1 (Quadratic program) Projection onto simplex.

Input: The function estimate components $\{f_p^{(n)}\}_{p=1}^P$, the data window $\mathcal{S}_L^{(n-1)}$ and the loss ℓ .

- 1: Set the components of \mathbf{a} and \mathbf{b} as in (5).
- 2: Sort \mathbf{b} in ascending order, denoted as \mathbf{u} : $u_1 \leq u_2 \leq \dots \leq u_P$ and rearrange \mathbf{a} accordingly into \mathbf{v} .
- 3: Find $\rho = \max \left\{ 1 \leq j \leq P : u_j - \frac{\sum_{i=1}^j \frac{u_i}{v_i} + 2}{\sum_{i=1}^j \frac{1}{v_i}} < 0 \right\}$.
- 4: Define $\mu = -\frac{2 + \sum_{i=1}^{\rho} \frac{u_i}{v_i}}{\sum_{i=1}^{\rho} \frac{1}{v_i}}$.
- 5: Compute $\theta_p = \max \left\{ -\frac{1}{2a_p} (b_p + \mu), 0 \right\}$ for $p = 1, \dots, P$.

Output: θ .

1) Every p th function estimate component $f_p^{(n)}$ is chosen through a single-kernel online method operating over the p th RK within the dictionary. The methods are selected by the user, and they can be different for each RK as long as all of them adopt the online setting described in Sec. III-A. For example, stochastic gradient descent methods for function estimation [35], and associated variants [36], [37], can be readily used. Since the function estimate components of $\mathbf{f}^{(n)}$ can be computed in parallel across P different computing units, the computational cost can be distributed.

2) Next, the convex weights in $\theta^{(n)}$ are chosen as the ones that minimize the partially evaluated upper bound cost (3c) at $\mathbf{f}^{(n)}$ and $\mathcal{S}_L^{(n-1)}$, referred from now on as learning cost. Mathematically,

$$\theta^{(n)} = \arg \min_{\theta \in \Delta^{P-1}} \check{\mathcal{C}}_{\eta} \left(\theta, \mathbf{f}^{(n)}; \mathcal{S}_L^{(n-1)} \right) \quad (4a)$$

$$= \arg \min_{\theta \in \Delta^{P-1}} \theta^{\top} \mathbf{A}^{(n)} \theta + \mathbf{b}^{(n)\top} \theta, \quad (4b)$$

where $\mathbf{A}^{(n)} \triangleq \text{diag}(\mathbf{a}^{(n)}) \in \mathbf{S}_{++}^P$, with $\mathbf{a}^{(n)} \in \mathbb{R}_+^P$, and $\mathbf{b}^{(n)} \in \mathbb{R}^P$ whose components are computed³ as

$$a_p^{(n)} = \frac{\eta}{2} \left\| f_p^{(n)} \right\|_{\mathcal{H}_p}^2, \quad (5a)$$

$$b_p^{(n)} = \sum_{i \in \mathcal{I}_L^{(n-1)}} \ell \left(f_p^{(n)}(x^{(i)}), y^{(i)} \right), \quad (5b)$$

for all $p \in \mathbb{N}^{[1,P]}$, where $\mathcal{I}_L^{(n-1)}$ corresponds to the index set associated with the data samples in $\mathcal{S}_L^{(n-1)}$. We adapt the projection onto the simplex algorithm discussed in [38], [39], [40] by extending its applicability to any quadratic problem described by a diagonal positive definite matrix with simplex constraints. As a result, the proposed **Algorithm 1** can solve (4) exactly. Its computational complexity is bottlenecked by a sorting step; that is, an asymptotic average complexity $\mathcal{O}(P \log P)$ [41]. It should be mentioned that this complexity can be further reduced to $\mathcal{O}(P)$ on average by using a randomized pivot algorithm variation that identifies the parameter ρ (**Algorithm 1**, line 3) using a divide and conquer procedure

³Notice that nothing prevents $f_p^{(n)}$ to be zero-valued and thus $\mathbf{A}^{(n)}$ from being singular and positive semi-definite. However, we can always set $\mathbf{A}^{(n)} = \text{diag}(\mathbf{a}^{(n)}) + \delta \mathbf{I}_P$ where δ is an arbitrarily small positive value.

instead of sorting [42], but this is out of the scope of the present paper.

3) Finally, the function estimate $f^{(n)} = \theta^{(n)\top} \mathbf{f}^{(n)}$ is proposed.

In summary, our scheme can be seen as a higher-level learner that iteratively chooses the lowest incurring learning cost combination of function estimates provided by lower-level learners, namely, single RK methods.

IV. PERFORMANCE ANALYSIS

Under an online setting, as the one described in Sec. III-A, the incurred cost accumulated over time receives the name of *cumulative* cost (CC). In our case, the CC up to the n th time step is given by $\sum_{i=1}^n \mathcal{C}_{\eta}(f^{(i)}; \mathcal{S}_L^{(i)})$. From here, recall that every i th function estimate $f^{(i)}$ is proposed via Sec. III-C, before the i th data window $\mathcal{S}_L^{(i)}$ becomes available; thus, the CC is a measure of performance protecting against overfitting. Intuitively, the lower the growth of the incurred CC with respect to n , the better the expected performance over unseen data. In fact, popular measures of performance, such as the dynamic regret, are constructed as the difference between the CC incurred by the sequence of function estimates proposed by a method and a sequence of comparators [43].

In order to validate our scheme experimentally, we pose a signal reconstruction online problem from synthetically generated streaming data. Specifically, we use the squared loss, i.e., $\ell(f(x), y) = (f(x) - y)^2$, and a dictionary of $P = 20$ Gaussian kernels, i.e., $k_p(x, t) = \exp(-\frac{1}{2}(x-t)^2/\sigma_p^2)$, with different widths σ_p linearly spaced between 0.1 and 10. The method associated with each RK is an augmented naive online R_{reg} minimization algorithm (NORMA) [35] with a window length of $L = 10$ data samples, a budget of $\tau = 100$ kernel expansion terms (beyond the allowed budget we truncate the oldest terms of the kernel expansion), and a fixed learning rate $\lambda_{\text{NORMA}} = 0.05$. That is, before any possible truncation, each n th function estimate associated with the p th RK is constructed as $f_p^{(n)} = \sum_{i=1}^{n-1} \alpha_{p,i}^{(n)} k_p(x^{(i)}, \cdot)$, where each $\alpha_{p,i}^{(n)} \in \mathbb{R}$ denotes a kernel-expansion coefficient obtained from the following NORMA update:

$$f_p^{(n)} = f_p^{(n-1)} - \lambda_{\text{NORMA}} \partial_f \mathcal{C}_{\eta} \left(f; \mathcal{S}_L^{(n-1)} \right) \Big|_{f=f_p^{(n-1)}}, \quad (6)$$

which, after some algebraic steps, leads to the next closed-form update rule:

$$\alpha_i^{(n)} = \begin{cases} -\lambda_{\text{NORMA}} \ell'_{p,i} \ell'^{(n-1)} & \text{if } i = n-1, \\ \gamma \alpha_i^{(n-1)} - \lambda_{\text{NORMA}} \ell'_{p,i} \ell'^{(n-1)} & \text{if } i \in \mathcal{I}_L^{(n-1)} \setminus \{n-1\}, \\ \gamma \alpha_i^{(n-1)} & \text{otherwise,} \end{cases} \quad (7)$$

where $\ell'_{p,i} \triangleq \ell' \left(f_p^{(n)}(x^{(i)}), y^{(i)} \right) = 2 \left(f_p^{(n)}(x^{(i)}) - y^{(i)} \right)$ and $\gamma \triangleq (1 - \lambda_{\text{NORMA}} \eta) \in \mathbb{R}_+^{(0,1)}$. The regularization parameter is chosen as $\eta = 0.01$. Lastly, the data samples have been generated via a stable AR(1) process $y^{(n)} = \varphi y^{(n-1)} + u^{(n)}$, with $\varphi = 0.5488135$, $u^{(n)} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 0.71519837)$, $y^{(0)} = 0$ and unit time stamps uniformly arranged in time, i.e., $x^{(n)} = n$.

Additionally, we compare our scheme with the online multiple kernel regression (OMKR) algorithm [21], arguably the closest approach conceptually. More specifically, we compare against the budget OMKR gradient-based variant method over the same experimental setting described above. Briefly, the considered OMKR method can be described, at each iteration step n , by the following three-stage scheme:

1) The set of function estimates, in this case, regressors proposed by each one of the P NORMAs, is updated as (7) and collected in $\mathbf{f}^{(n)} \in \mathcal{H}_{1:P}$.

2) Then, the P weights for combining the multiple regressors are updated as

$$\mathbf{w}^{(n)} = \mathbf{w}^{(n-1)} - \lambda_{\text{OMKR}}^{(n)} \nabla_{\mathbf{w}} \mathcal{C}_{\eta} \left(\mathbf{w}^{\top} \mathbf{f}^{(n)}; \mathcal{S}_L^{(n-1)} \right) \Big|_{\mathbf{w}=\mathbf{w}^{(n-1)}}. \quad (8)$$

In this case, we use an initial learning rate $\lambda_{\text{OMKR}}^{(1)} = 8 \cdot 10^{-4}$ that is halved every 50 steps until a minimum value of 10^{-5} . After some algebraic manipulations and making use of the definition of the RKHS direct sum norm [34], the evaluated gradient in (8) equals to

$$\sum_{i \in \mathcal{I}_L^{(n-1)}} \mathbf{f}^{(n)}(x^{(i)}) \ell' \left(\mathbf{w}^{(n-1)\top} \mathbf{f}^{(n)}(x^{(i)}), y^{(i)} \right) + \eta \mathbf{A}^{(n)} \mathbf{w}^{(n-1)}. \quad (9)$$

The matrix $\mathbf{A}^{(n)}$ corresponds to the one introduced in (5a), and the initial combination weights are set as $\mathbf{w}^{(1)} = \mathbf{0}_P$.

3) Finally, the function estimate $f^{(n)} = \mathbf{w}^{(n)\top} \mathbf{f}^{(n)}$ is proposed.

Unlike our scheme, the OMKR algorithm can eventually learn any linear combination of single-kernel function estimates. However, due to the additive nature of the update step in (8), the OMKR algorithm usually suffers from slow convergence rates. Moreover, it requires a sequence of learning rates $\lambda_{\text{OMKR}}^{(n)}$ whose tuning involves optimization techniques or task-specific knowledge, hence adversely affecting performance, e.g., poor learning or instabilities, if not carried out adequately.

Our experimental results in Fig. 2 show that the CC incurred by our proposed learning scheme outperforms the lowest CC incurred by any of the combined single RK NORMA regressors separately. In the same figure, it can also be observed that our scheme incurs a CC that increases at a lower rate than the one incurred by the OMKR algorithm, thus allowing our scheme to outperform the best NORMA regressor much sooner. The reason behind this observation is arguably the additive nature of the OMKR algorithm, which requires numerous updates to completely remove the residual contributions of irrelevant regressors. For example, see in Fig. 3 the “spiky” shape of the OMKR signal estimate due to some regressors constructed with a narrow-valued σ_p RK.

Regarding computational resources, both the OMKR algorithm and our scheme can be parallelized across the RKs within the dictionary, which in our experiments means a constant time complexity $\mathcal{O}(\tau)$, and a combination step of complexity $\mathcal{O}(P)$ and $\mathcal{O}(P \log P)$, respectively. However, as mentioned in Sec. III-C, the complexity of our scheme can

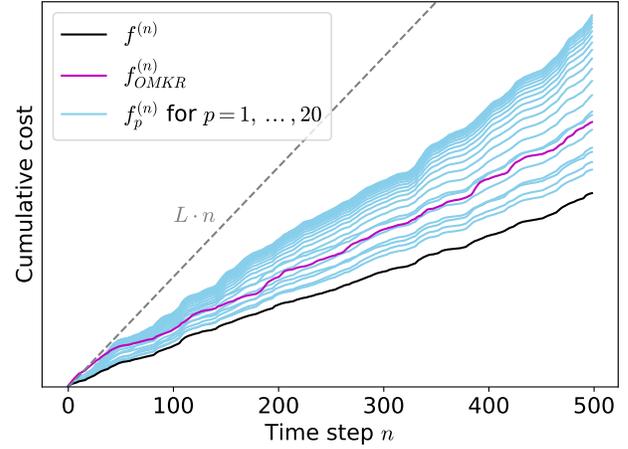


Fig. 2. Cumulative cost up to time step n incurred by our learning scheme, the OMKR algorithm, and the combined single RK NORMA regressors individually. The dashed-dotted line $L \cdot n$ is shown as a reference.

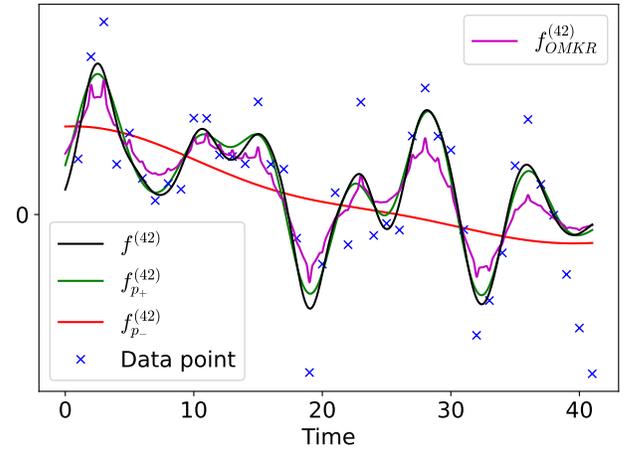


Fig. 3. Snapshot of the 42nd signal estimate obtained by the OMKR algorithm, our learning scheme, and the best and the worst of the combined single RK NORMA regressors (denoted by the indices p_+ and p_- , respectively) in terms of the so-far ($n = 42$) incurred cumulative cost.

be further reduced to $\mathcal{O}(P)$ on average, making both of the compared approaches computationally equivalent.

V. CONCLUSION

We present a multi-kernel learning scheme that experimentally outperforms the best of the combined single RK methods, in terms of the cumulative regularized least squares cost metric, with a comparable computational load per computing unit. This corroborates the ability of the proposed scheme to effectively accommodate a larger function space (from which to draw function estimates) of multi-kernel methods while keeping the lower computational complexity of online single RK methods. Furthermore, although **Algorithm 1** has been expressly designed for the task discussed in this paper, it can be used to solve any other problem that accepts a formulation as in (4b).

REFERENCES

- [1] B. Schölkopf and A. J. Smola, *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- [2] F. Pérez-Cruz and O. Bousquet, “Kernel methods and their potential use in signal processing,” *IEEE signal processing magazine*, vol. 21, no. 3, pp. 57–65, 2004.
- [3] H. Takeda, S. Farsiu, and P. Milanfar, “Kernel regression for image processing and reconstruction,” *IEEE Transactions on image processing*, vol. 16, no. 2, pp. 349–366, 2007.
- [4] A. Pozdnoukhov, *Prior Knowledge in Kernel Methods*. PhD thesis, Verlag nicht ermittelbar, 2006.
- [5] G. Wahba, *Spline models for observational data*. SIAM, 1990.
- [6] A. Nosedal-Sanchez, C. B. Storlie, T. C. Lee, and R. Christensen, “Reproducing kernel hilbert spaces for penalized regression: A tutorial,” *The American Statistician*, vol. 66, no. 1, pp. 50–60, 2012.
- [7] M. Stone, “Cross-validatory choice and assessment of statistical predictions,” *Journal of the royal statistical society: Series B (Methodological)*, vol. 36, no. 2, pp. 111–133, 1974.
- [8] O. Chapelle, V. Vapnik, O. Bousquet, and S. Mukherjee, “Choosing multiple parameters for support vector machines,” *Machine learning*, vol. 46, pp. 131–159, 2002.
- [9] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*, vol. 2. MIT press Cambridge, MA, 2006.
- [10] M. Gönen and E. Alpaydm, “Multiple kernel learning algorithms,” *The Journal of Machine Learning Research*, vol. 12, pp. 2211–2268, 2011.
- [11] F. Orabona and J. Luo, “Ultra-fast optimization algorithm for sparse multi kernel learning,” in *Proceedings of the 28th International Conference on Machine Learning*, 2011.
- [12] Z. Wang, K. Crammer, and S. Vucetic, “Breaking the curse of kernelization: Budgeted stochastic gradient descent for large-scale svm training,” *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 3103–3131, 2012.
- [13] J. A. Tropp, “Greed is good: Algorithmic results for sparse approximation,” *IEEE Transactions on Information theory*, vol. 50, no. 10, pp. 2231–2242, 2004.
- [14] Y. Engel, S. Mannor, and R. Meir, “The kernel recursive least-squares algorithm,” *IEEE Transactions on signal processing*, vol. 52, no. 8, pp. 2275–2285, 2004.
- [15] C. Richard, J. C. M. Bermudez, and P. Honeine, “Online prediction of time series data with kernels,” *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 1058–1067, 2008.
- [16] J. B. Souza Filho and P. S. Diniz, “A fixed-point online kernel principal component extraction algorithm,” *IEEE Transactions on Signal Processing*, vol. 65, no. 23, pp. 6244–6259, 2017.
- [17] J. B. Souza Filho, L.-D. Van, T.-P. Jung, P. S. Diniz, et al., “Online component analysis, architectures and applications,” *Foundations and Trends® in Signal Processing*, vol. 16, no. 3-4, pp. 224–429, 2022.
- [18] A. Rahimi and B. Recht, “Random features for large-scale kernel machines,” *Advances in neural information processing systems*, vol. 20, 2007.
- [19] R. Money, J. Krishnan, and B. Beferull-Lozano, “Sparse online learning with kernels using random features for estimating nonlinear dynamic graphs,” *IEEE Transactions on Signal Processing*, 2023.
- [20] O. Dekel, S. Shalev-Shwartz, and Y. Singer, “The forgetron: A kernel-based perceptron on a fixed budget,” *Advances in neural information processing systems*, vol. 18, 2005.
- [21] D. Sahoo, S. C. Hoi, and B. Li, “Online multiple kernel regression,” in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 293–302, 2014.
- [22] Y. Shen, T. Chen, and G. Giannakis, “Online ensemble multi-kernel learning adaptive to non-stationary and adversarial environments,” in *International Conference on Artificial Intelligence and Statistics*, pp. 2037–2046, PMLR, 2018.
- [23] D. Sahoo, S. C. Hoi, and B. Li, “Large scale online multiple kernel regression with application to time-series prediction,” *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 13, no. 1, pp. 1–33, 2019.
- [24] V. Vapnik, “Principles of risk minimization for learning theory,” *Advances in neural information processing systems*, vol. 4, 1991.
- [25] U. Marteau-Ferey, D. Ostrovskii, F. Bach, and A. Rudi, “Beyond least-squares: Fast rates for regularized empirical risk minimization through self-concordance,” in *Conference on learning theory*, pp. 2294–2340, PMLR, 2019.
- [26] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [27] C. A. Micchelli, M. Pontil, and P. Bartlett, “Learning the kernel function via regularization,” *Journal of machine learning research*, vol. 6, no. 7, 2005.
- [28] Y. Shen, T. Chen, and G. B. Giannakis, “Random feature-based online multi-kernel learning in environments with unknown dynamics,” *The Journal of Machine Learning Research*, vol. 20, no. 1, pp. 773–808, 2019.
- [29] J. Gorski, F. Pfeuffer, and K. Klamroth, “Biconvex sets and optimization with biconvex functions: a survey and extensions,” *Mathematical methods of operations research*, vol. 66, no. 3, pp. 373–407, 2007.
- [30] E. Hazan et al., “Introduction to online convex optimization,” *Foundations and Trends® in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.
- [31] F. Orabona, “A modern introduction to online learning,” *arXiv preprint arXiv:1912.13213*, 2019.
- [32] M. Zinkevich, “Online convex programming and generalized infinitesimal gradient ascent,” in *Proceedings of the 20th international conference on machine learning (icml-03)*, pp. 928–936, 2003.
- [33] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [34] N. Aronszajn, “Theory of reproducing kernels,” *Transactions of the American mathematical society*, vol. 68, no. 3, pp. 337–404, 1950.
- [35] J. Kivinen, A. J. Smola, and R. C. Williamson, “Online learning with kernels,” *IEEE transactions on signal processing*, vol. 52, no. 8, pp. 2165–2176, 2004.
- [36] A. Koppel, A. S. Bedi, K. Rajawat, and B. M. Sadler, “Optimally compressed nonparametric online learning: Tradeoffs between memory and consistency,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 61–70, 2020.
- [37] E. Ruiz-Moreno and B. Beferull-Lozano, “Tracking of quantized signals based on online kernel regression,” in *2021 IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6, IEEE, 2021.
- [38] S. Shalev-Shwartz and Y. Singer, “Efficient learning of label ranking by soft projections onto polyhedra,” *Journal of Machine Learning Research*, 2006.
- [39] W. Wang and M. A. Carreira-Perpinán, “Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application,” *arXiv preprint arXiv:1309.1541*, 2013.
- [40] M. Blondel, A. Fujino, and N. Ueda, “Large-scale multiclass support vector machine training via euclidean projection onto the simplex,” in *2014 22nd International Conference on Pattern Recognition*, pp. 1289–1294, IEEE, 2014.
- [41] A. D. Mishra and D. Garg, “Selection of best sorting algorithm,” *International Journal of intelligent information Processing*, vol. 2, no. 2, pp. 363–368, 2008.
- [42] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, “Efficient projections onto the l_1 -ball for learning in high dimensions,” in *Proceedings of the 25th international conference on Machine learning*, pp. 272–279, 2008.
- [43] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, “Online optimization: Competing with dynamic comparators,” in *Artificial Intelligence and Statistics*, pp. 398–406, PMLR, 2015.
- [44] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 1999.

Supplementary material for “An Online Multiple Kernel Parallelizable Learning Scheme”

CORRECTNESS OF ALGORITHM 1

A Lagrangian of problem (4) is

$$L(\boldsymbol{\theta}, \boldsymbol{\lambda}, \mu) = \boldsymbol{\theta}^\top \text{diag}(\mathbf{a}^{(n)})\boldsymbol{\theta} + \mathbf{b}^{(n)\top}\boldsymbol{\theta} - \boldsymbol{\lambda}^\top\boldsymbol{\theta} + \mu(\mathbf{1}^\top\boldsymbol{\theta} - 1), \quad (\text{S.1})$$

being $\mu \in \mathbb{R}$ and $\boldsymbol{\lambda} \in \mathbb{R}^P$ the Lagrange multipliers associated with the equality and inequality constraints, respectively. At the optimal solution $\boldsymbol{\theta}^{(n)}$, the following KKT conditions [44] hold:

$$2\theta_p^{(n)}a_p^{(n)} + b_p^{(n)} - \lambda_p + \mu = 0, \quad p = 1, \dots, P \quad (\text{S.2a})$$

$$\theta_p^{(n)} \geq 0, \quad p = 1, \dots, P \quad (\text{S.2b})$$

$$\lambda_p \geq 0, \quad p = 1, \dots, P \quad (\text{S.2c})$$

$$\lambda_p \theta_p^{(n)} = 0, \quad p = 1, \dots, P \quad (\text{S.2d})$$

$$\sum_{p=1}^P \theta_p^{(n)} = 1. \quad (\text{S.2e})$$

From the complementary slackness, stated in (S.2d), we can deduce that if the primal inequality constraint in (S.2b) is slacked, i.e., greater than zero, then $\lambda_p = 0$ and from the stationarity condition (S.2a), the solution fulfils

$$\theta_p^{(n)} = -\frac{1}{2a_p^{(n)}}(b_p^{(n)} + \mu) > 0. \quad (\text{S.3})$$

On the other hand, if the primal inequality constraint is tight, i.e., $\theta_p^{(n)} = 0$, then the dual constraint (S.2c) is not binding. Again, from the stationarity condition in (S.2a), we can identify those non-binding constraints as those that satisfy the following expression:

$$b_p^{(n)} + \mu = \lambda_p \geq 0. \quad (\text{S.4})$$

In this way, it is clear from (S.4) that the components of the optimal solution that are zero, if any, correspond to the larger components of $\mathbf{b}^{(n)}$. Without loss of generality, we can assume that the components of $\mathbf{b}^{(n)}$ are sorted in ascending order as long as the components of $\mathbf{a}^{(n)}$ are rearranged accordingly. Thus, by comparing $\mathbf{b}^{(n)}$ with the solution as follows:

$$\begin{aligned} b_1^{(n)} \leq \dots \leq b_\rho^{(n)} \leq b_{\rho+1}^{(n)} \leq \dots \leq b_P^{(n)}, \\ \theta_{\rho+1}^{(n)} = \dots = \theta_P^{(n)} = 0, \end{aligned} \quad (\text{S.5})$$

it can be concluded that the index $\rho \in \mathbb{N}^{[1,P]}$ determines the number of components in the solution that are nonzero. From here, and rewriting the equality primal constraint (S.2e) as

$$\sum_{p=1}^P \theta_p^{(n)} = \sum_{p=1}^{\rho} \theta_p^{(n)} = -\frac{1}{2} \sum_{p=1}^{\rho} \frac{1}{a_p^{(n)}} (b_p^{(n)} + \mu) = 1, \quad (\text{S.6})$$

the Lagrangian multiplier associated with the equality constraint can be isolated and computed as

$$\mu = -\frac{2 + \sum_{p=1}^{\rho} \frac{b_p^{(n)}}{a_p^{(n)}}}{\sum_{p=1}^{\rho} \frac{1}{a_p^{(n)}}}, \quad (\text{S.7})$$

as long as the index ρ is known.

Theorem 1: Let ρ be the number of positive components in the solution of optimization problem (4), then

$$\rho = \max \left\{ 1 \leq j \leq P : b_j - \frac{2 + \sum_{i=1}^j \frac{b_i}{a_i}}{\sum_{i=1}^j \frac{1}{a_i}} < 0 \right\}, \quad (\text{S.8})$$

where \mathbf{b} is obtained by sorting $\mathbf{b}^{(n)}$ components in ascending order and \mathbf{a} corresponds to $\mathbf{a}^{(n)}$ rearranged accordingly.

Proof. Let us first define the quantities $\varphi_j \triangleq b_j - (2 + \sum_{i=1}^j \frac{b_i}{a_i}) / \sum_{i=1}^j \frac{1}{a_i}$ and $s_{j:k} \triangleq \sum_{i=j}^k \frac{1}{a_i}$. Then, the goal is to show that $j = \rho$ is the largest index in $\{1, \dots, P\}$ for which φ_j remains negative.

For $j < \rho$, we have that

$$\varphi_j = \frac{1}{s_{1:j}} \left(s_{1:j} b_j - \left(2 + \sum_{i=1}^j \frac{b_i}{a_i} \right) \right) \quad (\text{S.9a})$$

$$= \frac{1}{s_{1:j}} \left(s_{1:j} b_j - 2 - \sum_{i=1}^{\rho} \frac{b_i}{a_i} + \sum_{i=j+1}^{\rho} \frac{b_i}{a_i} \right) \quad (\text{S.9b})$$

$$= \frac{1}{s_{1:j}} \left(s_{1:j} b_j + s_{1:\rho} \mu + \sum_{i=j+1}^{\rho} \frac{b_i}{a_i} \right) \quad (\text{S.9c})$$

$$= b_j + \mu + \frac{1}{s_{1:j}} \sum_{i=j+1}^{\rho} \frac{1}{a_i} (\mu + b_i) < 0, \quad (\text{S.9d})$$

where in step (S.9b) we use the equivalence $\sum_{i=1}^j \frac{b_i}{a_i} = \sum_{i=1}^{\rho} \frac{b_i}{a_i} - \sum_{i=j+1}^{\rho} \frac{b_i}{a_i}$. Next, in step (S.9c), we make use of the relation in (S.7). Finally, the step (S.9d) holds thanks to the relation in (S.3) and because $s_{1:j}, a_i \geq 0 \forall i, j$.

For $j = \rho$, and thanks to (S.3), we have $\varphi_{\rho} = b_{\rho} + \mu < 0$. Then, using the relation in (S.7), we can verify that (S.8) holds.

For $j > \rho$, we can follow similar algebraic steps as in (S.9) to obtain

$$\varphi_j = \frac{1}{s_{1:j}} \left(s_{1:\rho} (b_j + \mu) + \sum_{i=\rho+1}^j \frac{1}{a_i} (b_j - b_i) \right) \geq 0. \quad (\text{S.10})$$

The inequality in (S.10) holds thanks to the relation in (S.4) and the fact that $b_j \geq b_i \forall i$.